

CLAIMS

What is claimed is:

1. A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:
 - partitioning a file into a plurality of subfiles;
 - distributing the plurality of subfiles from a first node to a first group comprising a plurality of recipient nodes, wherein at least one subfile is distributed from the first node to each recipient node of said first group but no individual recipient node receives all of said plurality of subfiles; and
 - exchanging subfiles among said plurality of recipient nodes of said first group such that each recipient node of said first group obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile.
2. The method of claim 1 wherein said distributing comprising:
 - distributing from the first node a different subfile to each of said recipient nodes of said first group.
3. The method of claim 1 wherein said partitioning comprises:
 - partitioning said file into said plurality of subfiles corresponding in number to a number of said recipient nodes in said first group.
4. The method of claim 1 wherein said partitioning further comprises:
 - partitioning said file into said plurality of subfiles that are each approximately equal in size.
5. The method of claim 1 further comprising:
 - determining a number of said recipient nodes to include in said first group.

6. The method of claim 5 wherein said determining comprises:
determining a suitable number of concurrent communication connections that can be used for communication of information between one of the nodes and a plurality of the other nodes;
and
determining said number of recipient nodes to include in said first group as corresponding in number to said number of concurrent communication connections.
7. The method of claim 1 wherein said distributing comprises:
distributing the plurality of subfiles to said plurality of recipient nodes of said first group concurrently.
8. The method of claim 1 wherein exchanging subfiles among said plurality of recipient nodes of said first group further comprises:
each of said plurality of recipient nodes establishing concurrent communication connections to every other recipient node of said first group.
9. The method of claim 1 wherein said first node and said plurality of recipient nodes of said first group each comprise a server computer.
10. The method of claim 9 wherein said first node and said plurality of recipient nodes are distributed server computers in a Content Distribution Network (CDN).
11. The method of claim 1 further comprising:
said first group of recipient nodes communicating said file to a second group comprising a plurality of recipient nodes.
12. The method of claim 11 further comprising:
each recipient node of said first group communicating a subfile to every recipient node of said second group such that said recipient nodes of said second group each receive all of said plurality of subfiles.
13. The method of claim 11 further comprising:
each recipient node of said first group communicating the subfile that it receives from said first node to at least one node of the second group.

14. The method of claim 13 wherein each recipient node of said first group begins communicating the subfile that it is receiving from said first node to said at least one node of the second group before fully receiving the subfile from the first node.

15. The method of claim 1 further comprising:

logically organizing a plurality of groups of recipient nodes into a primary multicast tree, wherein the groups of the primary multicast tree are logically organized sequentially such that intermediate groups of the primary multicast tree each communicate the file to a next sequential group of the primary multicast tree and wherein each intermediate group begins to communicate the file to a next sequential group of the primary multicast tree before fully receiving the file from a preceding group of the primary multicast tree.

16. The method of claim 15 further comprising:

further logically organizing a plurality of groups of recipient nodes into a secondary multicast tree, wherein at least one group of the primary multicast tree begins communicating the file to at least one group of the secondary multicast tree after the group of the primary multicast tree has fully received the file.

17. A system comprising:

means for partitioning a file into a plurality of subfiles;

an origin node comprising means for distributing all of said plurality of subfiles from said origin node to a first group comprising a plurality of recipient nodes, wherein at least one subfile is distributed from the origin node to each recipient node of said first group but not all of said plurality of subfiles are distributed from the origin node to any of the recipient nodes of said first group; and

said recipient nodes of said first group each comprising means for exchanging their respective subfiles received from said origin node such that each recipient node of said first group obtains all of said plurality of subfiles, wherein said means for exchanging of at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node.

18. The system of claim 17 wherein said means for distributing comprises:

means for distributing from the origin node a different subfile to each of said recipient nodes of said first group.

19. The system of claim 17 wherein said means for partitioning comprises:
means for partitioning said file into said plurality of subfiles corresponding in number to a number of said recipient nodes in said first group.

20. The system of claim 17 wherein said means for distributing comprises:
means for distributing the plurality of subfiles to said plurality of recipient nodes of said first group concurrently.

21. The system of claim 17 wherein each of said recipient nodes of said first group further comprises:
means for communicating said file to a second group comprising a plurality of recipient nodes.

22. The system of claim 21 wherein said means for communicating of each recipient node of said first group begins communicating the subfile that such recipient node is receiving from said origin node to at least one node of the second group before fully receiving the subfile from the origin node.

23. A system comprising:
an origin node operable to partition a file into a plurality of subfiles, wherein said plurality of subfiles correspond in number to a number of recipient nodes in a first group to which said file is to be distributed;
said origin node operable to distribute all of said plurality of subfiles to said recipient nodes, wherein a different subfile is distributed from said origin node to each of said recipient nodes; and
said recipient nodes operable to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles, wherein at least one recipient node is operable to begin communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node before the at least one recipient node fully receives its respective subfile from the origin node.

24. The system of claim 23 wherein the origin node is operable to distribute the plurality of subfiles to said number of recipient nodes of said first group concurrently.

25. A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

distributing a plurality of subfiles that comprise a file from a first node to a first group comprising a plurality of recipient nodes, wherein at least one subfile distributed from the first node to each recipient node of said first group but no individual recipient node of said first group receives all of said plurality of subfiles from the first node; and

exchanging subfiles among said plurality of recipient nodes of said first group such that each recipient node of said first group obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile.

26. The method of claim 25 further comprising:
partitioning said file into said plurality of subfiles.

27. The method of claim 25 wherein said distributing comprising:
distributing from the first node a different subfile to each of said recipient nodes of said first group.

28. The method of claim 25 wherein said distributing comprises:
distributing the plurality of subfiles to said plurality of recipient nodes of said first group concurrently.

29. The method of claim 25 wherein exchanging subfiles among said plurality of recipient nodes of said first group further comprises:

each of said plurality of recipient nodes establishing concurrent communication connections to every other recipient node of said first group.

30. The method of claim 25 further comprising:
said first group of recipient nodes communicating said file to a second group comprising a plurality of recipient nodes.

31. The method of claim 30 wherein each recipient node of said first group begins communicating the subfile that it is receiving from said first node to at least one node of the second group before fully receiving the subfile from the first node.

32. A method of distributing from a first node to a plurality of recipient nodes a file encoded with multiple description coding, the method comprising:

distributing a plurality of descriptors of a file encoded with multiple description coding (MDC) from a first node to a first group comprising a plurality of recipient nodes, wherein at least one descriptor is distributed from the first node to each recipient node of said first group but not all of said plurality of descriptors are distributed from the first node to any of the recipient nodes of said first group; and

said plurality of recipient nodes of said first group exchanging their respective descriptors such that each recipient node of said first group obtains all of said plurality of descriptors, wherein at least one recipient node of said first group begins communicating a portion of its respective descriptor that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective descriptor from the first node.

33. The method of claim 32 wherein said distributing comprising:

distributing from the first node a different descriptor to each of said recipient nodes of said first group.

34. The method of claim 32 wherein said distributing comprises:

distributing the plurality of descriptors to said plurality of recipient nodes of said first group concurrently.